

What Is Claimed Is:

1. A method of buffering InfiniBand queue pairs in a single memory structure, comprising:
 - 5 receiving an InfiniBand packet comprising a portion of a communication to be transmitted on a non-InfiniBand communication link;
 - identifying a first queue pair to which the InfiniBand packet belongs;
 - identifying a first virtual lane to which said first queue pair belongs;
 - storing said portion of the communication in a first bucket of a single
 - 10 memory structure, wherein said single memory structure is configured to store contents of InfiniBand packets received from multiple queue pairs; and
 - for each of the multiple queue pairs, including said first queue pair, maintaining an associated linked list of buckets in said single memory structure in which contents of InfiniBand packets belonging to the queue pair are stored.
- 15 2. The method of claim 1, further comprising:
 - for each of the multiple queue pairs, maintaining:
 - a first pointer configured to identify the beginning of the associated linked list; and
 - 20 a second pointer configured to identify the end of the associated linked list.
- 25 3. The method of claim 1, further comprising:
 - maintaining a control structure comprising an entry corresponding to each bucket of said single memory structure;
 - wherein each entry in said control structure that is a member of a linked list associated with a queue pair is configured to identify a next control structure

entry and a next single memory structure bucket in said linked list.

4. The method of claim 3, further comprising:
5 updating a first entry in said control structure to reflect said storage of said portion of the communication.

5. The method of claim 1, further comprising:
10 updating an indicator configured to indicate a level of InfiniBand packets stored in said single memory structure for said first queue pair.

10 6. The method of claim 1, further comprising:
15 updating an indicator configured to indicate a level of InfiniBand packets stored in said single memory structure for said first virtual lane.

15 7. The method of claim 1, further comprising, prior to said storing:
20 determining whether sufficient space is available in said single memory structure to store said portion of the communication.

20 8. The method of claim 7, wherein said determining whether sufficient space is available comprises:
25 determining an amount of space in said single memory structure used to store portions of communications received via each queue pair belonging to said first virtual lane, including said first queue pair; and
comparing a sum of said determined amounts of space to an amount of space in said single memory structure allocated to said first virtual lane.

9. The method of claim 7, wherein said determining whether

sufficient space is available comprises:

determining an amount of space in said single memory structure used to store portions of communications received via said first queue pair; and

comparing said determined amount of space to an amount of space in said 5 single memory structure dedicated to said first queue pair.

10. The method of claim 9, wherein said determining whether sufficient space is available further comprises:

if said determined amount of space exceeds said dedicated amount of 10 space, determining whether a portion of said single memory structure used to store portions of communications received via multiple queue pairs has space available for storing said portion of the communication.

11. The method of claim 7, wherein said determining whether 15 sufficient space is available comprises:

determining a number of buckets in said single memory structure used to store portions of communications received via said first queue pair; and

comparing said number of buckets to a threshold number of buckets allocatable to said first queue pair.

20

12. The method of claim 1, wherein said single memory structure is a multi-port random access memory component.

13. The method of claim 1, wherein said control structure is a multi- 25 port random access memory component.

14. A computer readable medium storing instructions that, when

executed by a computer, cause the computer to perform a method of buffering InfiniBand queue pairs in a single memory structure, the method comprising:

receiving an InfiniBand packet comprising a portion of a communication to be transmitted on a non-InfiniBand communication link;

5 identifying a first queue pair to which the InfiniBand packet belongs;

identifying a first virtual lane to which said first queue pair belongs;

storing said portion of the communication in a first bucket of a single memory structure, wherein said single memory structure is configured to store contents of InfiniBand packets received from multiple queue pairs; and

10 for each of the multiple queue pairs, including said first queue pair,

maintaining an associated linked list of buckets in said single memory structure in which contents of InfiniBand packets belonging to the queue pair are stored.

15. The computer readable medium of claim 14, the method further comprising:

maintaining a control structure comprising an entry corresponding to each bucket of said single memory structure;

wherein each entry in said control structure that is a member of a linked list associated with a queue pair is configured to identify the next control structure entry and next single memory structure bucket in said linked list.

16. The computer readable medium of claim 14, the method further comprising:

25 updating an indicator configured to indicate a level of InfiniBand packets stored in said single memory structure for said first queue pair; and

updating an indicator configured to indicate a level of InfiniBand packets stored in said single memory structure for said first virtual lane.

17. In a communication device coupled to an InfiniBand network, a method of sharing one memory structure among multiple queue pairs, the method comprising:

5 receiving packets for each of multiple queue pairs terminating at the communication device; and

for each of said queue pairs:

reassembling in a shared memory contents of said packets into communications to be transmitted to an external communication system;

10 maintaining an associated linked list identifying locations in said shared memory in which said communications are reassembled; and tracking the amount of said shared memory being used to store contents of packets received via said queue pair.

15 18. The method of claim 17, further comprising:

for each of one or more virtual lanes, tracking the amount of said shared memory being used to store contents of packets received via said virtual lane.

19. The method of claim 17, wherein said reassembling comprises:

20 as said packets are received from the InfiniBand network, queuing said contents directly into said shared memory; wherein said shared memory serves as receive queues for each of said multiple queue pairs.

25 20. The method of claim 17, wherein said received packets comprise portions of encapsulated Ethernet packets.

21. The method of claim 17, wherein said maintaining an associated linked list for a first queue pair comprises:

5 maintaining a head pointer configured to identify:

a first location in said shared memory in which contents of a first packet are stored; and

10 a first entry in a shared control structure, said first entry corresponding to said first location in said shared memory; and maintaining a tail pointer configured to identify:

15 a final location in said shared memory in which contents of a final packet are stored; and

a final entry in said shared control structure;

wherein each entry in said shared control structure that is part of said first linked list, except for said final entry, identifies a subsequent entry in said shared control structure and identifies a location in said shared memory corresponding to 15 said subsequent entry.

22. The method of claim 17, further comprising:

managing said linked lists for said queue pairs with a shared control;

20 wherein each said location in said shared memory corresponds to an entry in said shared control; and

wherein each entry in said shared control is configured to identify:

a subsequent entry within the same linked list as said entry; and

25 a location in said shared memory corresponding to said subsequent entry.

23. A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of sharing one

memory structure among multiple queue pairs, the method comprising:

receiving packets for each of multiple queue pairs terminating at the communication device; and

for each of said queue pairs:

5 reassembling in a shared memory contents of said packets into communications to be transmitted to an external communication system; maintaining an associated linked list identifying locations in said shared memory in which said communications are reassembled; and tracking the amount of said shared memory being used to store
10 contents of packets received via said queue pair.

24. The computer readable medium of claim 23, the method further comprising:

for each of one or more virtual lanes, tracking the amount of said shared
15 memory being used to store contents of packets received via said virtual lane.

25. The computer readable medium of claim 23, wherein said reassembling comprises:

as said packets are received from the InfiniBand network, queuing said
20 contents directly into said shared memory;
 wherein said shared memory serves as receive queues for each of said multiple queue pairs.

26. The computer readable medium of claim 23, the method further comprising:

managing said linked lists for said queue pairs with a shared control;
 wherein each said location in said shared memory corresponds to an entry

in said shared control; and

wherein each entry in said shared control is configured to identify:

a subsequent entry within the same linked list as said entry; and

a location in said shared memory corresponding to said subsequent

5 entry.

27. A method of storing a communication received from an InfiniBand network, the method comprising:

receiving a set of InfiniBand packets from an InfiniBand network, each

10 said InfiniBand packet comprising a portion of a communication;

storing said communication portions in a memory shared among multiple queue pairs of the InfiniBand network, including a first queue pair through which said set of InfiniBand packets is received; and

maintaining a first linked list for said first queue pair to identify locations

15 in said memory in which said communication portions are stored.

28. The method of claim 27, wherein said storing comprises reassembling said communication portions into said communication.

20 29. The method of claim 27, wherein said maintaining a first linked list comprises:

in a control structure, maintaining a first linked list of control entries, wherein each of said control entries except a final control entry identifies:

a subsequent control entry; and

25 corresponding to said subsequent control entry, a location in said memory in which data received through said first queue pair are stored.

30. The method of claim 29, wherein said maintaining further comprises:

maintaining a head pointer identifying a first control entry in said first linked list and a first location in said memory; and

5 maintaining a tail pointer identifying said final control entry in said first linked list and a final location in said memory.

31. A computer readable medium storing instructions that, when executed by a computer, cause the computer to perform a method of storing a 10 communication received from an InfiniBand network, the method comprising:

receiving a set of InfiniBand packets from an InfiniBand network, each said InfiniBand packet comprising a portion of a communication;

storing said communication portions in a memory shared among multiple queue pairs of the InfiniBand network, including a first queue pair through which 15 said set of InfiniBand packets is received; and

maintaining a first linked list for said first queue pair to identify locations in said memory in which said communication portions are stored.

32. The computer readable medium of claim 31, wherein said 20 maintaining a first linked list comprises:

in a control structure, maintaining a first linked list of control entries, wherein each of said control entries except a final control entry identifies:

a subsequent control entry; and

25 corresponding to said subsequent control entry, a location in said memory in which data received through said first queue pair are stored.

33. An apparatus for storing contents of InfiniBand packets of one or

more communication streams, comprising:

 a receive module configured to receive InfiniBand packets from one or more communication streams;

 a single memory structure, wherein for each of the one or more

5 communication streams, buckets of said single memory structure in which contents of packets of the communication stream are stored are linked via a linked list associated with the communication stream; and

 a control structure configured to facilitate management of said linked list.

10 34. The apparatus of claim 33, wherein:

 said control structure comprises an entry corresponding to each bucket of said single memory structure; and

 each said entry in said control structure is configured to identify a next entry in said control structure and a corresponding next bucket in said single

15 memory structure.

35. The apparatus of claim 33, wherein each of said linked lists comprises:

20 the buckets of said single memory structure in which said contents of packets of the communication stream are stored; and

 for each said bucket in said single memory structure, a corresponding entry in said control structure.

36. The apparatus of claim 35, further comprising:

25 for each linked list associated with a communication stream:

 a first pointer identifying the beginning of said linked list; and
 a second pointer identifying the end of said linked list.

37. The apparatus of claim 33, wherein said single memory structure is a multi-port random access memory component.

5 38. The apparatus of claim 33, wherein said control structure is a multi-port random access memory component.

10 39. The apparatus of claim 33, wherein said control structure is configured to enable the one or more communication streams to make full use of said single memory structure.

40. The apparatus of claim 33, wherein the communication streams are virtual lanes.

15 41. The apparatus of claim 33, wherein the communication streams are queue pairs.